

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

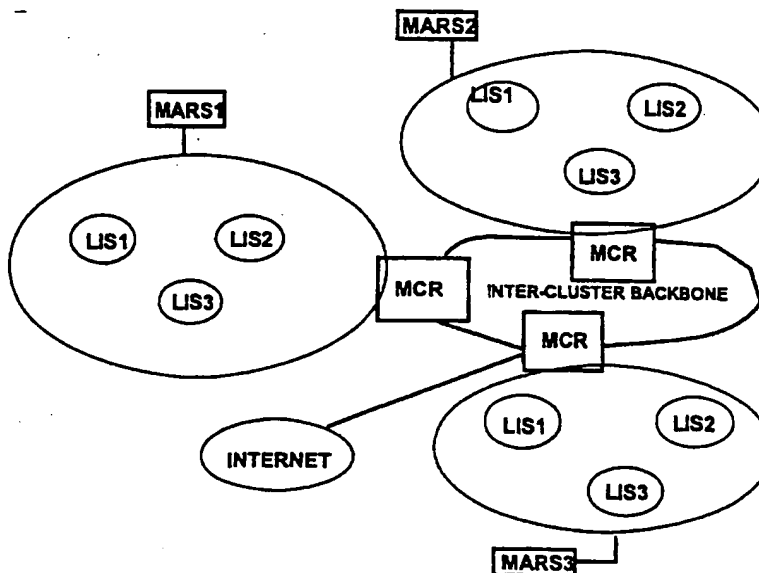
(51) International Patent Classification ⁶ : H04Q 11/04	A2	(11) International Publication Number: WO 99/29137 (43) International Publication Date: 10 June 1999 (10.06.99)
(21) International Application Number: PCT/SE98/02156 (22) International Filing Date: 27 November 1998 (27.11.98) (30) Priority Data: 9704457-2 1 December 1997 (01.12.97) SE (71) Applicant (for all designated States except US): TELIA AB (publ) [SE/SE]; Mårbackagatan 11, S-123 86 Farsta (SE). (72) Inventor; and (75) Inventor/Applicant (for US only): KAVAK, Nail [SE/SE]; Myrstuguvägen 359, S-143 32 Värby (SE). (74) Agent: PRAGSTEN, Rolf; Telia Research AB, Vitsandsgatan 9, S-123 86 Farsta (SE).		(81) Designated States: EE, JP, LT, LV, NO, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published Without international search report and to be republished upon receipt of that report.

(54) Title: IP/ATM NETWORK SYSTEM ADAPTED FOR THE SIMULTANEOUS TRANSMISSION OF IP DATA PACKETS TO A PLURALITY OF USERS

(57) Abstract

The invention provides an ATM transmission system which is adapted for the simultaneous transmission (i.e. multicasting) of IP data packets to a plurality of system users, using ATM as a carrier network, and which includes, an IP/ATM network including clusters of Logical IP Systems (LISs), a multicast address resolution server (MARS) for each cluster, and a multicast router (MCR) for each cluster. The clusters are interconnected through the MCRs which form an inter-cluster backbone for the IP/ATM network. Each cluster includes an integral number of ATM hosts and LISs, the size of a cluster is dependent on physical and administrative constraints, and each LIS, within a cluster, is adapted to support a number of ATM hosts. Multicasting of IP data packets within a cluster

utilizes a MARS protocol, and multicasting between clusters utilizes an inter-cluster multicasting protocol (ILMP). All ATM hosts, within a cluster, even if supported by different LISs, are adapted to be configured with an address of the same MARS. The system includes means for gaining access to the Internet, for example, a non-ATM IP Subnet connected to said inter-cluster backbone. Alternatively, one of the MCRs may be adapted for use as an entry point for the IP/ATM network to the Internet.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

- 1 -

IP/ATM NETWORK SYSTEM ADAPTED FOR THE SIMULTANEOUS TRANSMISSION OF IP DATA PACKETS TO A PLURALITY OF USERS

5 The present invention relates to an ATM transmission system which utilizes an inter-cluster multicast protocol (ILMP), a method for the simultaneous transmission (i.e. multicasting) of IP data packets using ATM as a carrier network, and a telecommunication system including the ATM transmission system of the present invention.

10 In multicasting, one message is simultaneously transmitted to a number of receivers, i.e. transmitted at the same time to a number of users of a multicasting service. With a multicasting service, use of the network's resources may be optimized by copying the multicast messages in the net rather than at the user terminals. Many of the existing services, offered by network operators, could have been implemented, in a more effective manner, if support for multicasting had been
15 available for use with such services. In future, many multimedia applications will be dependent on a multicasting service. Some examples of multimedia applications, which could be more effectively implemented with the aide of a multicasting service, are as follows:

- 20 - exchange of information in real time;
- video conferences;
- distance working, for example, working at a location remote from a
25 central office, company premises, administration centre, etc.; and
- video games.

30 IP (Internet Protocol) networks, in which ATM (Asynchronous Transfer Mode) is used as a carrier network, do not have support for multicasting. Known existing solutions, for example, MARS (Multicast Address Resolution Service), work for a limited number of users, i.e. about 255 users.

- 2 -

The Internet Engineering Task Force (IETF) have published a Request for Comment (RFC), namely, RFC 2022, which specifies a method that provides multicast support for IP end-systems, or routers, over an ATM network. However, the RFC 2022 (MARS) solution is restricted to intra-LIS communication.

5

In Telia's IP/ATM network, known as Broadway, multicasting can be provided using the MARS protocol. However, such a solution does not scale since any multicasting protocol has to deal with both intra-LIS and inter-LIS multicasting.

10

There is, therefore, at the present time, a distinct lack of established methods for inter-LIS multicasting for use in IP/ATM telecommunication systems.

15

It is an object of the present invention to provide an ATM transmission system in which multicasting between clusters utilizes inter-cluster multicasting protocols, and in which multicasting within a cluster utilizes MARS. A cluster is a collection of hosts and routers belonging to the same domain. The use of an inter-cluster multicasting protocol (ILMP) makes it possible to provide a multicast service to a much larger number of users, i.e. up to 65,000 users, than is possible using present day solutions, for example, MARS which is, in essence, limited to one LIS and about 255 users.

20

It is another object of the present invention to provide a method for the simultaneous transmission (i.e. multicasting) of IP data packets, using ATM as a carrier network.

25

It is a further object of the present invention to provide a telecommunication system including the ATM transmission system of the present invention, or operating in accordance with the method of the present invention.

30

According to one aspect of the present invention, there is provided, an ATM transmission system, adapted for the simultaneous transmission (i.e. multicasting) of IP data packets to a plurality of system users, using ATM as a carrier network, characterised in that said transmission system includes an IP/ATM network including

- 3 -

clusters of Logical IP Systems (LISs); a multicast address resolution server (MARS) for each cluster; and a multicast router (MCR) for each cluster, in that said clusters are interconnected through said MCRs which form an inter-cluster backbone for said IP/ATM network, in that multicasting of IP data packets within a cluster utilizes a MARS protocol, and in that multicasting between clusters utilizes an inter-cluster multicasting protocol (ILMP). The ATM transmission system may include means for gaining access to the Internet. Said means for gaining access to the Internet may include a non-ATM IP Subnet connected to said inter-cluster backbone. One of said MCRs may be adapted for use as an entry point for the IP/ATM network to the Internet.

Each cluster may consist of an integral number of ATM hosts and LISs, the size of a cluster being dependent on physical and administrative constraints, each LIS, within a cluster, may be adapted to support a number of ATM hosts, and all ATM hosts within a cluster, even if supported by different LISs, may be adapted to be configured with an address of the same MARS. For each cluster, the functionality of - MARS and MCR may be adapted to be co-located.

- In the presence of multiple MCRs belonging to the same cluster, said IP/ATM network may be adapted to select a MCR with the highest IP address as a default MCR.

MARS may be adapted to operate in server mode.

Said MCRs may be adapted to promiscuously receive all multicast traffic on a cluster.

Said MCRs may be adapted to support point-to-multipoint connections provided by UNI 3.0 signalling.

In accordance with the ATM transmission system of the present invention, all ATM hosts and MCRs, associated with the same cluster, may be adapted to be

- 4 -

configured with the address of the same MARS, MARS may be adapted to operate in server mode, said IP/ATM network may be adapted to enable IP/ATM endpoints, within a cluster, to register as a member of an inter-cluster multicast group utilizing MARS procedures, said MCRs may be adapted to support point-to-multipoint connections provided by UNI 3.0 signalling, and a multicast server may be provided and adapted to manage said point-to-multipoint connections, and to distribute multicast traffic therefrom (i.e. as root) to all other leaves of said multipoint connections.

10 A MCR may be adapted, in advance of multicasting, to select a well known end-node and establish a connection thereto, said well known end-node being adapted to provide a database for the ATM addresses of inter-cluster multicast group members, and said well known end-node may be adapted to establish a point-to-multipoint tree and to distributing multicast traffic sent by a member of said inter-cluster multicast group to other members of said group using said tree. Each group address may have a corresponding well known end-node address.

15 Said IP/ATM network may be adapted to enable multicast packets to be sent to members of an inter-cluster multicast group, even if the sender is not a member of said group.

20 Said MCRs may be adapted to operate in either a server mode, or a Virtual Circuit (VC) mesh mode, the mode used being dependent on several criteria, such as, the number of inter-cluster multicast groups, the grouping mode (i.e. sparse, dense, etc.), traffic density, charging policies of a network operator, performance, reliability and system complexity.

25 When MCRs is adapted to operate in a server mode, said well-known end-node may be adapted to provide a database for IP multicast addresses, and corresponding ATM addresses for members of an inter-cluster multicast group, said MCRs may be adapted to issue multicast packets to said well known end-node for distribution to members of an inter-cluster multicast group, and said well known end-

- 5 -

node may be adapted to establish a point-to-multipoint tree to said group members and to distributing multicast packets, received from said MCRs, to said group members using said tree.

5 When MCRs is adapted to operate in VC mesh mode, said well known end-nodes may be adapted to keep track of IP multicast addresses, and corresponding ATM members of said MCRs, and said IP/ATM network may be adapted to report changes in membership of an MCR to a respective well known end-node, and to
10 establish point-to-multipoint trees between MCRs belonging to the same inter-cluster multicast group(s), when an association between said IP multicast addresses and said ATM members has been determined. Said IP/ATM network may include two control VCs between a MCR and said well-known end-node, in which case, a first one of said control VCs is a transient bi-directional point-to-point connection, between a
15 MCR and said well-known end-node, and is adapted for use in registering group membership, and a second one of said control VCs is a unidirectional point-to-multipoint group_ctrl_VC, from said well-known end-node to the MCR, and is adapted to propagate group joining and leaving messages to the MCR.

 Said IP/ATM network may be adapted to enable a multicast packet originating
20 from a non-group member to be distributed to group members using a point-to-multipoint tree of a group member MCR, said well known end-node being adapted to establish a point-to-point connection between said non-group member and said group member MCR to facilitate the transfer of said multicast packet to said group member MCR and thereafter to said group members via said point-to-multipoint tree.

25 Said IP/ATM network may be adapted to enable a MCR to join, or leave, an inter-cluster multicast group, at the option of the MCR.

30 According to another aspect of the present invention, there is provided, in an ATM transmission system including an IP/ATM network, said network including clusters of Logical IP Systems (LISs), a multicast address resolution server (MARS) for each cluster and a multicast router (MCR) for each cluster, a method for the

- 6 -

simultaneous transmission (i.e. multicasting) of IP data packets to a plurality of system users using ATM as a carrier network, said method being characterised by interconnecting said clusters through said MCRs to form an inter-cluster backbone for said IP/ATM network, by using a MARS protocol for multicasting said IP data packets within a cluster, and by using an inter-cluster multicasting protocol (ILMP) for multicasting said IP data packets between clusters. The method may be characterised by obtaining access to the Internet for said IP/ATM network. The method may be further characterised by said IP/ATM network, in the presence of multiple MCRs belonging to the same cluster, selecting a MCR with the highest IP address as a default MCR. MARS may operate in server mode. Said MCRs may promiscuously receiving all multicast traffic on a cluster and may support point-to-multipoint connections provided by UNI 3.0 signalling.

The method may be characterised by configuring all ATM hosts and MCRs, associated with the same cluster, with the address of the same MARS; operating MARS in server mode; IP/ATM endpoints, within a cluster, utilizing MARS procedures to register as a member of an inter-cluster multicast group; said MCRs supporting point-to-multipoint connections provided by UNI 3.0 signalling; and using a multicast server to manage said point-to-multipoint connections and distribute multicast traffic therefrom (i.e. as root) to all other leaves of said multipoint connections.

The method may be characterised by a MCR, in advance of multicasting selecting a well known end-node, and establishing a connection to said well known end-node, said well known end-node providing a database for ATM addresses of inter-cluster multicast group members; and by said well known end-node establishing a point-to-multipoint tree to said ATM addresses and, on receipt of a multicast packet from a member of said inter-cluster multicast group, using said tree to distribute said multicast packet to other members of said group. Each group address may have a corresponding well known end-node address.

The method may be characterised by said well known end-node forming the root of a point-to-multipoint tree when said MCRs are operating in the server mode;

- 7 -

and by an MCR, wishing to distribute a multicast packet to members of a existing group of which it is not a member, establishing a point-to-point connection to said well-known end-node, and sending said multicast packet to said well-known end-node using the point-to-point connection; and, on receipt of said multicast packet, said well known end-node distributing said multicast packet, via an existing point-to-multipoint tree to said group members.

The method may be characterised by said MCRs operating in either a server mode, or a Virtual Circuit (VC) mesh mode, the mode used being dependent on several criteria, such as, the number of inter-cluster multicast groups, the grouping mode (i.e. sparse, dense, etc.), traffic density, charging policies of a network operator, performance, reliability and system complexity. The method may be further characterised by said MCRs operating in a server mode; said well-known end-node providing a database for IP multicast addresses, and corresponding ATM addresses for members of an inter-cluster multicast group; said MCRs issuing multicast packets to said well known end-node for distribution to members of an inter-cluster multicast group; and said well known end-node establishing a point-to-multipoint tree to said group members and using said tree to distribute multicast packets, received from said MCRs, to said group members. The method may be further characterised by said MCRs operating in VC mesh mode; said well known end-nodes keeping track of IP multicast addresses, and corresponding ATM members of said MCRs; and said IP/ATM network reporting changes in membership of an MCR to a respective well known end-node; and, on determining an association between said IP multicast addresses and said ATM members, establishing point-to-multipoint trees between MCRs belonging to the same inter-cluster multicast group(s). The method may be further characterised by said IP/ATM network establishing two control VCs between a MCR and said well-known end-node, a first one of said control VCs being a transient bi-directional point-to-point connection, between a MCR and said well-known end-node, and a second one of said control VCs being a unidirectional point-to-multipoint group_ctrl_VC, from said well-known end-node to the MCR; using said first control VC to register membership of a an inter-cluster multicast group; and using said second control VC to propagate group joining and leaving messages to the MCR.

- 8 -

The method may be characterised by said MCRs, at start-up, establishing a point-to-point connection with a well-known end-node; and informing said end-node of their ATM addresses and of the inter-cluster multicast groups from which they want to receive multicast traffic; on receipt of information concerning the multicast group(s) from which an MCR wants to receive multicast traffic, said well known end-node returning to said MCR, the ATM addresses which support said groups; and said MCR, on receipt of said ATM addresses, establishing point-to-multipoint connections with each of the multicast group members. The method may be further characterised by a MCR, on receipt of a multicast packet from its cluster, using said point-to-multipoint tree to forward the multicast packet to other clusters. The method may be further characterised by said IP/ATM network using a point-to-multipoint tree of a group member MCR to distribute a multicast packet originating from a non-group member to group members; and said well known end-node establishing a point-to-point connection between said non-group member and said group member MCR to facilitate transfer of said multicast packet to said group member MCR and thereafter to said group members via said point-to-multipoint tree. The method may be further characterised by a MCR, wishing to leave an inter-cluster multicast group, contacting said well-known end-node; said well known end-node, on receipt of said leaving request, propagating a leave message, in respect of said MCR, to the other MCRs on its group membership list; and each of said other MCRs, on receipt of said leave message, dropping a leaf for said leaving MCR from their respective point-to-multipoint trees. The method may be further characterised by a MCR, wishing to join an inter-cluster multicast group, contacting said well-known end-node; said well known end-node, on receipt of the joining request, returning the ATM addresses of existing group members to said new MCR and propagating the ATM address of said new MCR to existing group members, through a point-to-multipoint tree; said new MCR member, on receipt of the ATM addresses, establishing a point-to-multipoint connection to each of said existing group members; and each of said existing group members, on receipt of the ATM address of the new MCR, updating their respective point-to-multipoint trees to include a leaf for said new MCR.

The method may be characterised by said IP/ATM network operating in server

- 9 -

mode ILMP, in conjunction with MARS; MCRs, on obtaining information concerning existing groups, propagating said information to a corresponding well known end-node, residing within an inter-cluster backbone, any one of said MCRs acting as said well known end-node for a multicast group, each group having only one well known end-node; a MCR sending group membership reports to corresponding well-known end-nodes via a transient control channel established between said MCR and the well known end-node, said control channel being released after a predetermined period of time; a MCRs establishing point-to-point data channels for tunneling multicast packets in unicast fashion and a point-to-multipoint tree for each multicast group using itself as the root of the tree; existing group members reporting their own ATM addresses to said well-known end-node when a MCR propagates information concerning the existing groups on their cluster side; and using group member IDs to eliminate multicast packet reflections, said IDs being assigned by said well known end-node. The method may be further characterised by MCRs selecting a well known end-node for a particular inter-cluster multicast group by invoking algorithmic mapping of multicast addresses and corresponding well known end-nodes.

The method may be characterised by an ATM host, wishing to transmit a multicast packet to members of a multicast group, sending said packet to a Multicast Server (MCS) entity, on receipt of said packet, said MCS coping the packet to all its interfaces where there are group members within the cluster, including the MCR (since MCRs behave promiscuously) and unicasting said multicast packet to a corresponding well known end-node via a preestablished data channel; and, if said multicast group is reported to said well-known end-node, said well known end-node extracting the multicast packet from the unicast packet and sending the extracted multicast packet to group members using its point-multipoint tree to said members. The method may be further characterised by a MCS wishing to participate in a multicast group, configuring two addresses, namely, ATM addresses of primary and secondary (back-up) well known end-nodes. The method may be further characterised by a MCR maintaining a single bi-directional-control connection to the well-known end-node for registration/de-registration purposes; a single data unidirectional data channel for tunneling multicast packets; and a number (N) of

- 10 -

incoming multicast VCs, where N is the total number of multicast groups.

5 The method may be characterised by, in the event of failure of a primary well known end-node, dismantling all point-to-point connections between the MCS routers and said primary well-known end-node; the MCRs, on detecting dismantlement of said connections, opening a new circuit to a back-up, pre-configured, well known end-node; each MCR registering the ATM addresses of the group members of which it is aware to the back-up well known end-node; said back-up well known end-node, on receipt of the ATM addresses, establishing point-multipoint trees for each group; 10 said MCRs continuously try to reestablish a connection with said primary well known end-node; and on reestablishment of said connection, said MCRs dismantle their connection to said back-up well known end-node and re-establish their connection to said primary well known end-node.

15 According to a further aspect of the present invention, there is provided, a telecommunication system, characterised in that said system includes an ATM transmission system as outlined in preceding paragraphs, or is adapted to operate in accordance with the method outlined in preceding paragraphs.

20 The foregoing and other features of the present invention will be better understood from the following description with reference to the accompanying drawings, in which:

25 Figure 1 diagrammatically illustrates, in the form of a block diagram, an architecture for an inter-cluster multicasting network;

 Figure 2 diagrammatically illustrates, in form of a block diagram, a Virtual Circuit (VC) mesh for a number of MCRs;

30 Figure 3 diagrammatically illustrates, in the form of a block diagram, the VC mesh of Figure 2 of the drawing, including an MCR for a non-group member sending a packet to a multicast group;

- 11 -

Figure 4 diagrammatically illustrates, in the form of a block diagram, a point-to-multipoint tree constructed from a known node to all members of a multicast group; and

5 Figure 5 diagrammatically illustrates, in the form of a block diagram, an example scenario for a network architecture adapted to provide inter-cluster communication, where each cluster consists of multiple LISs.

10 In order to facilitate an understanding of the present invention a glossary of terms used in this patent specification is provided below:

	ATM:	Asynchronous Transfer Mode
	CMP:	Cluster Multicasting Protocol
15	IETF:	Internet Engineering Task Force
	ILMP:	Inter-cluster Multicast Protocol
20	IP:	Internet Protocol
	LIS:	Logical IP System/Subnet
	MARS:	Multicast Address Resolution Service
25	MCR:	Multicast Router
	MCS:	Multicast Server
30	RFC:	Request for Comment
	UNI:	User-Network Interface

- 12 -

VC: Virtual Circuit

VCI: Virtual Channel Identifier

5 VPI: Virtual Path Identifier

It will be seen from the subsequent description of the present invention that:

- 10 - the invention is concerned with inter-cluster multicasting;
- multicasting within a cluster utilizes IETF MARS;
- a cluster is a collection of ATM hosts and routers that belong to the same domain;
- 15 - a cluster may comprise a plurality of Logical IP Systems (LISs); and
- the size of a cluster depends on physical and administrative constraints, but, when choosing the size of a cluster, implementation is a very important
- 20 consideration.

In essence, an ATM transmission system of the present invention is adapted for the simultaneous transmission (i.e. multicasting) of IP data packets to a plurality of system users, using ATM as a carrier network, and includes:

- 25 - an IP/ATM network including clusters of Logical IP Systems (LISs), a cluster being adapted to support a number of ATM hosts and being of a size that is dependent on physical and administrative constraints;
- 30 - a multicast address resolution server (MARS) for each cluster, all ATM hosts within a cluster, even if supported by different LISs, are adapted to be configured with an address of the same MARS; and

- 13 -

- a multicast router (MCR) for each cluster, the functionality of MCR and MARS may be co-located for each cluster..

5 The clusters are interconnected through the MCRs which form an inter-cluster backbone for the IP/ATM network and the multicasting of IP data packets within a cluster utilizes a MARS protocol, whilst multicasting between clusters utilizes an inter-cluster multicasting protocol (ILMP). The ATM transmission system includes means for gaining access to the Internet, for example, a non-ATM IP Subnet connected to
10 the inter-cluster backbone. Alternatively, one of the MCRs may be adapted for use as an entry point for the IP/ATM network to the Internet.

 It is not recommended to have a very large number of end-systems in a cluster because this may cause performance degradation and difficulty in managing very
15 large groups. Intra-cluster communication between group members occurs directly (shortcut), whereas inter-cluster communication is effected through a router having a multicasting capability. The use of multicast routers as an aggregation point for inter-cluster communication helps to reduce the number of virtual connections, as well as the signalling load on the end-systems.

20 The architecture for an inter-cluster multicasting network of the present invention, which is diagrammatically illustrated, in the form of a block diagram, in Figure 1 of the accompanying drawings, includes three clusters, each one of which includes three Logical IP Systems, LIS₁, LIS₂ and LIS₃. As illustrated in Figure 1,
25 each cluster has an address resolution server associated therewith, namely, MARS₁, MARS₂ and MARS₃, and a multicast router (MCR). The three clusters are interconnected through the multicast routers (MCRs) which form an inter-cluster backbone for the network. One of the MCRs is used as an entry point for the network to the Internet.

30

 Intra-cluster communication is largely based on the MARS architecture, the only deviation being that the current MARS architecture assumes that each cluster

- 14 -

corresponds to only one LIS whereas, with the arrangement of Figure 1, each cluster encompass multiple LISs. It will be seen from the subsequent description that mapping multiple LISs to one cluster can be readily achieved. MARS is fully documented in RFC 2022, published by IETF, and its details will not, therefore, be described in this patent specification.

The assumptions that have been made in relation to the architecture of the inter-cluster multicasting network of Figure 1 are that:

- (a) The architecture uses, by and large, MARS protocol for intra-cluster multicasting and ILMP for inter-cluster multicasting;
- (b) A cluster consists of an integral number of hosts and LISs. For each cluster there is one address resolution server (MARS) and one multicast router (MCR). Although the functionality of MARS and MCR can, as previously stated, be co-located.
- (c) For scalability and reliability reasons it is possible to physically distribute the functionality of MARS within a cluster.
- (d) All ATM hosts residing on the same cluster, even if they belong to different LISs, are configured with the address of the same MARS.
- (e) MARS operates in server mode. Virtual Circuit (VC) mesh mode, as such, is not addressed in this patent specification because it is not necessary for an understanding of the present invention.
- (f) In the presence of multiple MCRs belonging to the same cluster, the MCR with the highest IP address is chosen as the default MCR.
- (g) MCRs promiscuously receive all multicast traffic on a Cluster.

- 15 -

- (h) The multicast routers (MCRs) support point-to-multipoint (pt-to-mpt) connections provided by UNI 3.0 signalling only.
- 5 (i) There is a corresponding well known end-node address for each group address.
- (j) Well known routers are selected algorithmically, as described in the subsequent description of the present invention.
- 10 (k) A sender can send packets to a multicast group even if the sender is not a member of the group.

The manner in which the inter-cluster multicasting network of Figure 1 operates will now be described. Within a cluster, IP/ATM endpoints that would like to join a multicasting group, use the MARS procedures to register. All ATM hosts (and
15 multicast routers), belonging to the same cluster, are configured with the address of the same MARS. A multicast server entity (assuming that MARS operates in server mode) manages the point-to-multipoint connections, and distributes the multicast traffic from itself (root) to all other leaves.

20

As illustrated in Figure 1, the MCRs inter-connect, or act as aggregation points between, clusters. In operation, an MCR selects a well known end-node with which it establishes communication, in advance of the multicasting process. The procedures used to select a well known end-node are outlined in the subsequent
25 description. The function of the well known end-node is that it serves as a database for the ATM addresses of the members of an inter-cluster multicast group. The well known end-node is adapted to establish a point-to-multipoint tree and to distributing multicast traffic sent by a member of the inter-cluster multicast group to other members of the multicast group using the tree, i.e. the end-node acts as the root of
30 the tree and is adapted to send IP multicast data packets, across the cluster boundaries, to the leaves of the tree, each one of which is the ATM address of an existing group member. Thus, in the server mode, the selected well known end-node

- 16 -

distributes the multicast traffic, sent by any member of the multicast group, to other members of the group via a previously constructed point-to-multipoint tree.

5 A MCR is adapted to operate in two different modes, namely a Virtual Circuit (VC) mesh mode and a server mode. Both modes of operation have their advantages and disadvantages and the operation mode which is chosen depends on several criteria, such as:

- 10 - the number of inter-cluster multicast groups;
- the grouping mode, i.e. sparse, dense, etc.;
- traffic density;
- 15 - charging policies of the network operators;
- performance;
- reliability;
- 20 - complexity, etc..

In the VC mesh mode, the well-known end-nodes are adapted to keep track of the IP multicast addresses, and the corresponding ATM members of the MCRs
25 [(IP1: ATM1, ATM2,...ATMN), (IP2: ATM1, ATM2,...ATMN),.....(IPN: ATM1, ATM2,...ATMN)]. Each and every change in membership of an MCR is reported to the well-known end-node. Once the association between the IP multicast addresses and ATM members is determined, point-to-multipoint trees are established between each MCR belonging to the same inter-cluster multicast group(s). A Virtual Circuit
30 (VC) mesh for a number of MCRs is diagrammatically illustrates, in form of a block diagram, in Figure 2 of the accompanying drawings.

- 17 -

In the server mode, the well-known end-node, in addition to being a database for the IP multicast addresses, and the corresponding ATM addresses, is adapted to relay multicast packets to members of an inter-cluster multicast group. MCRs are adapted to issue multicast packets to the well-known end-node, in unicast fashion, for distribution to members of the inter-cluster multicast group. A well known end-node is adapted to establish a point-to-multipoint tree to members of an associated inter-cluster multicast group, and to distribute multicast packets, received from a MCR, to group members using the tree. In other words, a point-to-multipoint tree is constructed from the well-known end-node to each of the group members, and multicast packets, received from a MCR, are forwarded to the group members via the point-to-multipoint (pt-to-mpt) tree.

In the VC mesh mode (or source routed point-to-multipoint), there are two control VCs between a MCR and the well-known end-node, which are as follows:

- (1) a transient bi-directional point-to-point connection, between a MCR and the well-known end-node, to register group membership; and
- (2) a unidirectional pt-to-mpt group_ctrl_VC, from the well-known end-node back to the MCR, to propagate group join and leave messages.

At start-up, the MCRs establish a point-to-point connection with the well-known end-node and inform the well known end-node of their ATM addresses, as well as the multicast groups from whom they want to receive traffic. The well known end-node, in turn, returns the ATM addresses to the MCRs which support the same groups. Once the MCRs receive the ATM addresses of the other group members, they open up point-to-multipoint connections with each of the multicast MCR group members. At the end of this process, initialization will have been effected and a bidirectional tree will have been constructed for every group.

30

When a MCR promiscuously receives a multicast packet from its cluster, it uses the previously constructed point-to-multipoint tree, as described above, to

- 18 -

forward the multicast packet to other clusters, as is diagrammatically illustrated in Figure 2 of the accompanying drawings, i.e. every group member manages a point-to-point tree.

5 In those situations, as is diagrammatically illustrated in Figure 3 of the accompanying drawings, where a non-group member, i.e. a sender (ATM host) within a particular cluster having a Multicast Router, MCR_4 , wishes to send a multicast packet to an inter-cluster multicast group which has members in other clusters to which the sender belongs, the Multicast Router, MCR_4 , issues a query to the well-known end-node to find out if the group exists in other clusters. The well-known end-node, on receipt of this query, returns the ATM address of an "appropriate" group member to which it establishes a point-to-point connection. The non-member MCR_4 then tunnels the multicast packet, in unicast fashion, to the "appropriate" group member (root). As is diagrammatically illustrated in Figure 3, the "appropriate" group member (root) is MCR_1 , which, on receipt of the multicast packet from MCR_4 , forwards the packet through its point-to-multipoint connections to the other Multicast Routers, MCR_2 and MCR_3 .

20 A Multicast Router (MCR), wishing to leave a group, simply contacts the well-known end-node which, in turn, propagates a leave message, in respect of that MCR, to the other MCRs on its group membership list. On receipt of the leave message, each of the MCRs drop the leaf for the leaving MCR from their point-to-multipoint trees.

25 A Multicast Router (MCR), wishing to join an existing multicast group, simply contacts the well-known end-node. On receipt of the request to join the group, the well-known end-node returns the ATM addresses of the existing group members, and the new MCR member, on receipt of the addresses, establishes a point-to-multipoint connection to each of the existing group members. At the same time, the well-known end-node propagates the ATM address of the newly joined MCR to the other group members, through the point-to-multipoint tree, to enable the existing members to update their trees to include a leaf for the new MCR member.

30

- 19 -

As is diagrammatically illustrated in Figure 4 of the accompanying drawings, the well known end-node forms the root of the point-to-multipoint tree when the MCRs are operating in the server mode. In other words, in server mode, the well know end-node, on receipt of the list of groups, establishes a point-to-multipoint tree back to the group members, MCR₁, MCR₂ and MCR₃. In this case, if there is a new incoming MCR member to an existing MCR group, a membership notification is sent to the well-known end-node which simply adds the new MCR member to its point-to-multipoint tree. When a MCR member wishes to terminate its membership of an existing group, the MCR issues a notification to the well-known end-node which drops the leaf for that MCR from its point-to-multipoint tree.

MCRs, wishing to distribute multicast packets, who are not members of any existing group can simply establish a connection to the well-known end-node which will thereafter distribute the multicast packet, via the point-to-multipoint tree managed by the well know end node.

The detailed operation of a server mode ILMP, in conjunction with MARS, will now be described. End-systems, within a cluster, use MARS procedures (MARS-REQUEST, MARS_MULTI, MARS_NAK, etc.) to register their ATM addresses and the list of the addresses of the groups they wish to join. The MARS entity, operating in the server mode, will itself establish point-to-multipoint tree prior to packet transmission. Information concerning the detailed operation of MARS is given in IETF RFC 2022.

It is assumed that an MCR obtains knowledge of existing groups (not members) within a cluster. This is readily achievable when MARS and the MCR are co-located. Thus, the detailed operation of the server mode ILMP, in conjunction with MARS, will be as follows:

- (1) Once the MCR obtains information concerning existing groups, it propagates this information to a corresponding well known end-node that resides within the inter-cluster backbone (see Figure 1 of the accompanying drawings).

- 20 -

Membership reports from MCRs to the well-known end-nodes are sent through a transient control channel. This control channel is released after a predetermined time.

- 5 (2) The MCRs also establish a point-to-point data channel for tunneling multicast packets in unicast fashion. Finally, the MCRs establish a point-to-multipoint tree for every group from itself back to its leaves, i.e. MCRs.
- 10 (3) Any of the MCRs can act as the well-known end-node but for a given group there is only one well-known end-node. When a MCR propagates the existing groups on their cluster side, the existing group members report their own ATM addresses to the well-known end-node.
- 15 (4) A 16 bit extended encapsulation carrying a 16-bit cluster member ID is used (as proposed in MARS) to solve the packet reflection problem. This ID is assigned by the well-known end-node.

Multicast forwarding is effected as follows:

- 20 (1) When a host wishes to transmit a multicast packet, it sends the packet to a Multicast Server (MCS) entity, which copies the packet to all its interfaces where there are group members within the cluster, including the MCR (since MCRs behave promiscuously).
- 25 (2) The multicast packet is unicast to the well known end-node via the preestablished data channel.
- 30 (3) If a multicast group is reported to the well-known end-node, after extracting the multicast packet from the unicast packet, the well known end-node will simply issue the packet through its point-multipoint tree.

A MCS wishing to participate in a multicast group, only needs to configure two

- 21 -

addresses, namely, the ATM addresses of the primary and the secondary (back-up) well known end-nodes.

5 Thus, the total number of circuits that will be maintained by any router will be as follows:

- a single bi-directional-control connection to the well-known end- node for registration/de-registration purposes;
- 10 - a single data unidirectional data channel for tunneling multicast packets; and
- a number (N) of incoming multicast VCs, where N is the total number of multicast groups.

15

The manner in which a well known end-node is selected will now be described. It is assumed that, within the multicast backbone, there are a number of well-known end-nodes each serving for a multicast group prefix. Mapping of multicast addresses and corresponding well-known end-nodes can be performed algorithmically. Given a group address and a set-of well-known end-nodes, the
20 algorithm will return the well-known end-node to be used for a group. Disregarding where the algorithm is executed, the same well-known end-node will always be returned for a given group. In practice, the routers at the edge of the multicast backbone, i.e. the MCRs, invoke the algorithmic mapping to select a well-known end-
25 node for a particular group.

In order to ensure that the well known end-node is not a single point of failure, a back-up node is provided which is adapted, in the event of failure of a well known end-node, to take over from the failed well known end-node.

30

Thus, if a well known end-node fails, for whatever reason:

- 22 -

- all point-to-point connections between the MCS routers and the well-known end-node are dismantled;
- 5 - the routers, on detecting that the connections have been dismantled, are adapted to open a new circuit to a back-up (pre-configured) well known end-node, assuming that the ATM address of the back-up well known end-node is also known a priori;
- 10 - thereafter, each of the MCRs register the ATM addresses of the group members of which it is aware to the back-up well known end-node; and
- 15 - on receipt of the ATM addresses, the back-up well known end-node establishes point-multipoint trees for each group.

It should be noted that, in general, multicast servers constitute a single point of failure and, in this respect, the server mode is weaker than the VC mesh mode.

20 In practice, routers will continuously try to establish a connection with the primary well known end-node. Whenever this succeeds, the routers are adapted to dismantle their connection to the back-up well known end-node and re-establish their connection to the primary end-node.

25 Consider an example scenario for a network architecture adapted to provide inter-cluster communication, i.e. adapted to operate in server mode ILMP, in conjunction with MARS, in which there are two multicast groups, each including a number of ATM hosts, namely, group G1 and group G2, which are as follows:

- 30 - G1: {ATM₁, ATM₂, ATM₃, ATM₄, ATM₅}; and
- G2: {ATM₁, ATM₂, ATM₁₀};

- 23 -

and three clusters:

- CLUSTER_x: ATM₁, ATM₂, ATM₃, ATM₄;
- 5 - CLUSTER_y: ATM₅, ATM₆, ATM₇; and
- CLUSTER_z: ATM₈, ATM₉, ATM₁₀

Figure 5 of the accompanying drawings diagrammatically illustrates such a
 10 scenario, in the form of a block diagram, in which the three clusters, CLUSTER_x,
 CLUSTER_y and CLUSTER_z, each include a number of LISs, which each have a
 number of ATM hosts connected thereto. In this example, CLUSTER_x includes three
 LISs, LIS₁ to LIS₃, CLUSTER_y includes two LISs, LIS₄ and LIS₅, and CLUSTER_z
 15 includes one LIS, LIS₆. The LISs of CLUSTER_x are connected to a ROUTER_x, the
 LISs of CLUSTER_y are connected to a ROUTER_y and the LIS of CLUSTER_z is
 connected to a ROUTER_z via a MCR.

Each Router (MARS plus MCR) is connected to an Inter-cluster Multicast
 Backbone, as is a non-ATM IP Subnet for providing entry to the Internet.

20

All IP hosts set up a bi-directional control channel to their respective MCRs
 (ROUTERS_{x,z}) for registration/de-registration purposes. Hosts register their ATM
 addresses, as well as the IP multicast groups they wish to be a member of through
 a respective channel. This channel is dismantled if not used for a predetermined
 25 period of time. ROUTERS_{x,z} keep track of {layer 3, ATM₁, ATM₂, ATM₃,ATM_N}
 mappings. For each and every registration request, the ROUTERS_{x,z} open a point-to-
 multipoint tree to all their respective members.

At the end of this phase:

30

- ROUTER_x will have an entry in its mapping table similar to:

{G1: ATM₁, ATM₂, ATM₃, ATM₄, ROUTER_x}, {G2: ATM₁, ATM₂, ROUTER_x};

- 24 -

- ROUTER_y will have an entry in its mapping table similar to:

[[G1: ATM₅, ROUTER_y]]; and

5

- ROUTER_z will have an entry in its mapping table similar to:

[[ATM₁₀, ROUTER_z]].

10 The ROUTERS_{xz} will also propagate existing groups, within the clusters, to a wide-area in the following manner:

- (1) A Router sets up a point-to-point connection to a well-known end-node and informs all multicast groups of which it wants to be a member, of its ATM address; and
- (2) The well known end-node, on receipt of this information, sets up a point-to-multipoint tree for each group.

15

20 According to the example given above, there will be one point-to-multipoint tree for G1, starting from a well-known end-node to ROUTER_x and ROUTER_y, and one point-to-multipoint tree for G2, starting from a well-known end-node to ROUTER_x and ROUTER_z.

25 When an ATM host, for example, ATM₁, wishes to send a multicast packet to group G1:

- the multicast packet is sent to a MCS within the cluster, i.e. CLUSTER_x;
- 30 - the MCS copies the packet to its internal point-to-multipoint channel (VPI/VCI_x) leading to hosts ATM₁, ATM₂, ATM₃, ATM₄, as well as to the well known end-node;

- 25 -

- the well known end-node, on receipt of the copied multicast packet, distributes the packet to ROUTER_x and ROUTER_y;
- 5 - the ROUTER_y distribute the incoming multicast packet to its internal interfaces, i.e. ATM_s, which resides within its own CLUSTER_y; and
- the packet that will be reflected to ROUTER_x is discarded using the extended encapsulation mechanism, referred to in a preceding paragraph.

10

The procedure to be followed when a host wishes to leave a group will now be described, by way of example, for the case where host, ATM₁, from CLUSTER_x and host, ATM₁₀, from CLUSTER_z decide to leave the group G2:

- 15 - the host, ATM₁, issues a leave request to ROUTER_x and host, ATM₁₀, issues a leave request to ROUTER_z;
- on receipt of the leave request, ROUTER_x and ROUTER_z respectively delete the leaves for hosts, ATM₁ and ATM₁₀, from their internal tree;
- 20 - at the same time, ROUTER_z detects that there are no group members left in CLUSTER_z for group G2 and sends a leave request to the well known end-node; and
- 25 - the well known end-node, on receipt of the leave request, removes ROUTER_z from its tree for G2.

It will be seen from the foregoing description that multicasting inside a cluster uses the IETF MARS (Multicast Address Resolution Service) protocol and that, in
30 accordance with the present invention, an inter-cluster multicasting protocol (CMP) is used for the multicasting of IP data packets between clusters. The CMP protocol implies that there are gateways between clusters to spread the multicast address to

- 26 -

other gateways and that the multicast traffic flow is such that multicast messages, leaving a cluster, are only transmitted to gateways where multicast (group) members exist. Each group has a predetermined unit, i.e. a well known end-node, with which respective gateways are associated. After registering group addresses to chosen well known end-nodes, if a gateway sees that there are members spread to other clusters they will set up point-to-multipoint ATM connections between gateways. Each gateway belonging to the same multicast group will later set up a point-to-multipoint connection to the rest of the members. Multicast packages, transmitted to a group, are not copied between clusters if there are no members therein, i.e. the packages are copied locally. To increase reliability and redundancy, a back-up node is used. If the node goes down, the gateways set up a new ATM connection between themselves from the information they have collected from the first registration.

Each gateway is configured with a number of IP addresses particular to a number of predetermined nodes. There may also be a number of predetermined nodes where each node serves a segment of existing multicast group addresses. The mapping of a multicast address and the associated predetermined node is algorithmic.

The present invention is primarily intended to be used in IP/ATM exchanges/gateways and the use of an inter-cluster multicasting protocol will enable a multicast message service to be provided to up to 65,000 users.

25

30

CLAIMS

1. An ATM transmission system, adapted for the simultaneous transmission (i.e.
5 multicasting) of IP data packets to a plurality of system users, using ATM as a carrier network, characterised in that said transmission system includes an IP/ATM network including clusters of Logical IP Systems (LISs); a multicast address resolution server (MARS) for each cluster; and a multicast router (MCR) for each cluster, in that said clusters are interconnected through said MCRs which form an inter-cluster backbone
10 for said IP/ATM network, in that multicasting of IP data packets within a cluster utilizes a MARS protocol, and in that multicasting between clusters utilizes an inter-cluster multicasting protocol (ILMP).
2. An ATM transmission system, as claimed in claim 1, characterised in that said
15 system includes means for gaining access to the Internet.
3. An ATM transmission system, as claimed in claim 2, characterised in that one
20 of said MCRs is adapted for use as an entry point for the IP/ATM network to the Internet.
4. An ATM transmission system, as claimed in claim 2, characterised in that said
means for gaining access to the Internet includes a non-ATM IP Subnet connected
to said inter-cluster backbone.
- 25 5. An ATM transmission system, as claimed in any preceding claim, characterised in that each cluster consists of an integral number of ATM hosts and LISs, the size of a cluster being dependent on physical and administrative constraints, and in that each LIS, within a cluster, is adapted to support a number of ATM hosts, and in that all ATM hosts within a cluster, even if supported by different LISs, are
30 adapted to be configured with an address of the same MARS.
6. An ATM transmission system, as claimed in any preceding claim,

- 28 -

characterised in that, for each cluster, the functionality of MARS and MCR is adapted to be co-located.

- 5 7. An ATM transmission system, as claimed in any preceding claim, characterised in that, in the presence of multiple MCRs belonging to the same cluster, said IP/ATM network is adapted to select a MCR with the highest IP address as a default MCR.
- 10 8. An ATM transmission system, as claimed in any preceding claim, characterised in that MARS is adapted to operate in server mode.
- 15 9. An ATM transmission system, as claimed in any preceding claim, characterised in that said MCRs are adapted to promiscuously receive all multicast traffic on a cluster.
- 20 10. An ATM transmission system, as claimed in any preceding claim, characterised in that said MCRs are adapted to support point-to-multipoint connections provided by UNI 3.0 signalling.
- 25 11. An ATM transmission system, as claimed in any preceding claim, characterised in that all ATM hosts and MCRs, associated with the same cluster, are adapted to be configured with the address of the same MARS, in that MARS is adapted to operate in server mode, in that said IP/ATM network is adapted to enable IP/ATM endpoints, within a cluster, to register as a member of an inter-cluster multicast group utilizing MARS procedures, in that said MCRs are adapted to support point-to-multipoint connections provided by UNI 3.0 signalling, and in that a multicast server is provided and adapted to manage said point-to-multipoint connections, and to distribute multicast traffic therefrom (i.e. as root) to all other leaves of said multipoint connections.
- 30 12. An ATM transmission system, as claimed in any preceding claim, characterised in that a MCR is adapted, in advance of multicasting, to select a well

- 29 -

known end-node and establish a connection thereto, said well known end-node being adapted to provide a database for the ATM addresses of inter-cluster multicast group members, and in that said well known end-node is adapted to establish a point-to-multipoint tree and to distributing multicast traffic sent by a member of said inter-cluster multicast group to other members of said group using said tree.

13. An ATM transmission system, as claimed in claim 12, characterised in that each group address has a corresponding well known end-node address.

10 14. An ATM transmission system, as claimed in any preceding claim, characterised in that said IP/ATM network is adapted to enable multicast packets to be sent to members of an inter-cluster multicast group, even if the sender is not a member of said group.

15 15. An ATM transmission system, as claimed in any preceding claim, characterised in that said MCRs are adapted to operate in either a server mode, or a Virtual Circuit (VC) mesh mode, the mode used being dependent on several criteria, such as, the number of inter-cluster multicast groups, the grouping mode (i.e. sparse, dense, etc.), traffic density, charging policies of a network operator, performance,
20 reliability and system complexity.

16. An ATM transmission system, as claimed in claim 15, when appended to any of claims 12 to 14, characterised in that said MCRs are adapted to operate in a server mode, in that said well-known end-node is adapted to provide a database for IP
25 multicast addresses, and corresponding ATM addresses for members of an inter-cluster multicast group, in that said MCRs are adapted to issue multicast packets to said well known end-node for distribution to members of an inter-cluster multicast group, and in that said well known end-node is adapted to establish a point-to-multipoint tree to said group members and to distributing multicast packets, received
30 from said MCRs, to said group members using said tree.

17. An ATM transmission system, as claimed in claim 15, when appended to any

- 30 -

5 of claims 12 to 14, characterised in that said MCRs are adapted to operate in VC mesh mode, in that said well known end-nodes are adapted to keep track of IP multicast addresses, and corresponding ATM members of said MCRs, and in that said IP/ATM network is adapted to report changes in membership of an MCR to a respective well known end-node, and to establish point-to-multipoint trees between MCRs belonging to the same inter-cluster multicast group(s), when an association between said IP multicast addresses and said ATM members has been determined.

10 18. An ATM transmission system, as claimed in claim 17, characterised in that said IP/ATM network includes two control VCs between a MCR and said well-known end-node; in that a first one of said control VCs is a transient bi-directional point-to-point connection, between a MCR and said well-known end-node, and is adapted for use in registering group membership, and in that a second one of said control VCs is a unidirectional point-to-multipoint group_ctrl_VC, from said well-known end-node
15 to the MCR, and is adapted to propagate group joining and leaving messages to the MCR.

20 19. An ATM transmission system, as claimed in any of claims 12 to 18, characterised in that said IP/ATM network is adapted to enable a multicast packet originating from a non-group member to be distributed to group members using a point-to-multipoint tree of a group member MCR, said well known end-node being adapted to establish a point-to-point connection between said non-group member and said group member MCR to facilitate the transfer of said multicast packet to said group member MCR and thereafter to said group members via said point-to-multipoint
25 tree.

30 20. An ATM transmission system, as claimed in any preceding claim, characterised in that said IP/ATM network is adapted to enable a MCR to join, or leave, an inter-cluster multicast group, at the option of the MCR.

21. In an ATM transmission system including an IP/ATM network, said network including clusters of Logical IP Systems (LISs), a multicast address resolution server

- 31 -

(MARS) for each cluster and a multicast router (MCR) for each cluster, a method for the simultaneous transmission (i.e. multicasting) of IP data packets to a plurality of system users using ATM as a carrier network, said method being characterised by interconnecting said clusters through said MCRs to form an inter-cluster backbone for said IP/ATM network, by using a MARS protocol for multicasting said IP data packets within a cluster, and by using an inter-cluster multicasting protocol (ILMP) for multicasting said IP data packets between clusters.

22. A method, as claimed in claim 24, characterised by obtaining access to the Internet for said IP/ATM network.

23. A method, as claimed in claim 21, or claim 22, characterised by said IP/ATM network, in the presence of multiple MCRs belonging to the same cluster, selecting a MCR with the highest IP address as a default MCR.

24. A method, as claimed in any of claims 21 to 23, characterised by MARS operating in server mode.

25. A method, as claimed in any of claims 21 to 24, characterised by said MCRs promiscuously receiving all multicast traffic on a cluster.

26. A method, as claimed in any of claims 21 to 25, characterised by said MCRs supporting point-to-multipoint connections provided by UNI 3.0 signalling.

27. A method, as claimed in any of claims 21 to 26, characterised by:

- configuring all ATM hosts and MCRs, associated with the same cluster, with the address of the same MARS;
- operating MARS in server mode;
- IP/ATM endpoints, within a cluster, utilizing MARS procedures to register as

- 32 -

a member of an inter-cluster multicast group;

- said MCRs supporting point-to-multipoint connections provided by UNI 3.0 signalling; and

5

- using a multicast server to:

- manage said point-to-multipoint connections; and

10

- distribute multicast traffic therefrom (i.e. as root) to all other leaves of said multipoint connections.

28. A method, as claimed in any of claims 21 to 29, characterised by:

15

- a MCR, in advance of multicasting:

- selecting a well known end-node; and

20

- establishing a connection to said well known end-node, said well known end-node providing a database for ATM addresses of inter-cluster multicast group members; and

- said well known end-node:

25

- establishing a point-to-multipoint tree to said ATM addresses; and

- on receipt of a multicast packet from a member of said inter-cluster multicast group, using said tree to distribute said multicast packet to other members of said group.

30

29. A method, as claimed in claim 28, characterised by each group address having a corresponding well known end-node address.

- 33 -

30. A method, as claimed in any of claim 21 to 29, characterised by:

- 5 - said well known end-node forming the root of a point-to-multipoint tree when said MCRs are operating in the server mode; and
- an MCR, wishing to distribute a multicast packet to members of a existing group of which it is not a member.
- 10 - establishing a point-to-point connection to said well-known end-node; and
- sending said multicast packet to said well-known end-node using the point-to-point connection; and
- 15 - on receipt of said multicast packet, said well known end-node distributing said multicast packet, via an existing point-to-multipoint tree to said group members.

20 31. A method, as claimed in any of claims 21 to 30, characterised by said MCRs operating in either a server mode, or a Virtual Circuit (VC) mesh mode, the mode used being dependent on several criteria, such as, the number of inter-cluster multicast groups, the grouping mode (i.e. sparse, dense, etc.), traffic density, charging policies of a network operator, performance, reliability and system

25 complexity.

32. A method, as claimed in claim 31, when appended to any of claims 28 to 30, characterised by:

- 30 - said MCRs operating in a server mode;
- said well-known end-node providing a database for IP multicast addresses,

- 34 -

and corresponding ATM addresses for members of an inter-cluster multicast group;

5 - said MCRs issuing multicast packets to said well known end-node for distribution to members of an inter-cluster multicast group; and

10 - said well known end-node:

- establishing a point-to-multipoint tree to said group members; and
- using said tree to distribute multicast packets, received from said MCRs, to said group members.

15 33. A method, as claimed in claim 31, when appended to any of claims 28 to 30, characterised by:

20 - said MCRs operating in VC mesh mode;

20 - said well known end-nodes keeping track of IP multicast addresses, and corresponding ATM members of said MCRs; and

25 - said IP/ATM network:

- reporting changes in membership of an MCR to a respective well known end-node; and
- on determining an association between said IP multicast addresses and said ATM members, establishing point-to-multipoint trees between MCRs belonging to the same inter-cluster multicast group(s).

30

34. A method, as claimed in claim 33, characterised by said IP/ATM network:

- 35 -

- 5 - establishing two control VCs between a MCR and said well-known end-node, a first one of said control VCs being a transient bi-directional point-to-point connection, between a MCR and said well-known end-node, and a second one of said control VCs being a unidirectional point-to-multipoint group_ctrl_VC, from said well-known end-node to the MCR;
- 10 - using said first control VC to register membership of a an inter-cluster multicast group; and
- 10 - using said second control VC to propagate group joining and leaving messages to the MCR.
- 35. A method, as claimed in any of claims 32 to 34, characterised by:
 - 15 - said MCRs, at start-up:
 - establishing a point-to-point connection with a well-known end-node; and
 - 20 - informing said end-node of their ATM addresses and of the inter-cluster multicast groups from which they want to receive multicast traffic;
 - 25 - on receipt of information concerning the multicast group(s) from which an MCR wants to receive multicast traffic, said well known end-node returning to said MCR, the ATM addresses which support said groups; and
 - 30 - said MCR, on receipt of said ATM addresses, establishing point-to-multipoint connections with each of the multicast group members.
- 36. A method, as claimed in claim 35, characterised by a MCR, on receipt of a multicast packet from its cluster, using said point-to-multipoint tree to forward the

- 36 -

multicast packet to other clusters.

37. A method, as claimed in claim 35, or claim 36, characterised by:

- 5 - said IP/ATM network using a point-to-multipoint tree of a group member MCR to distribute a multicast packet originating from a non-group member to group members; and
- 10 - said well known end-node establishing a point-to-point connection between said non-group member and said group member MCR to facilitate transfer of said multicast packet to said group member MCR and thereafter to said group members via said point-to-multipoint tree.

38. A method, as claimed in claim 35, or claim 36, characterised by:

- 15 - a MCR, wishing to leave an inter-cluster multicast group, contacting said well-known end-node;
- 20 - said well known end-node, on receipt of said leaving request, propagating a leave message, in respect of said MCR, to the other MCRs on its group membership list; and
- 25 - each of said other MCRs, on receipt of said leave message, dropping a leaf for said leaving MCR from their respective point-to-multipoint trees.

39. A method, as claimed in claim 35, or claim 36, characterised by:

- 30 - a MCR, wishing to join an inter-cluster multicast group, contacting said well-known end-node;
- said well known end-node, on receipt of the joining request:

- 37 -

- returning the ATM addresses of existing group members to said new MCR; and
 - propagating the ATM address of said new MCR to existing group members, through a point-to-multipoint tree;
- 5
- said new MCR member, on receipt of the ATM addresses, establishing a point-to-multipoint connection to each of said existing group members; and
- 10
- each of said existing group members, on receipt of the ATM address of the new MCR, updating their respective point-to-multipoint trees to include a leaf for said new MCR.
40. A method, as claimed in any of claims 31 to 39, characterised by:
- 15
- said IP/ATM network operating in server mode ILMP, in conjunction with MARS;
- 20
- MCRs, on obtaining information concerning existing groups, propagating said information to a corresponding well known end-node, residing within an inter-cluster backbone, any one of said MCRs acting as said well known end-node for a multicast group, each group having only one well known end-node;
- 25
- a MCR sending group membership reports to corresponding well-known end-nodes via a transient control channel established between said MCR and the well known end-node, said control channel being released after a predetermined period of time;
- 30
- a MCRs establishing:
 - point-to-point data channels for tunneling multicast packets in unicast fashion; and

- 38 -

- a point-to-multipoint tree for each multicast group using itself as the root of the tree;
 - 5 - existing group members reporting their own ATM addresses to said well-known end-node when a MCR propagates information concerning the existing groups on their cluster side; and
 - 10 - using group member IDs to eliminate multicast packet reflections, said IDs being assigned by said well known end-node.
41. A method, as claimed in claim 40, characterised by MCRs selecting a well known end-node for a particular inter-cluster multicast group by invoking algorithmic mapping of multicast addresses and corresponding well known end-nodes.
- 15
42. A method, as claimed in claim 40, or claim 41, characterised by:
- an ATM host, wishing to transmit a multicast packet to members of a multicast group, sending said packet to a Multicast Server (MCS) entity;
 - 20 - on receipt of said packet, said MCS:
 - coping the packet to all its interfaces where there are group members within the cluster, including the MCR (since MCRs behave promiscuously); and
 - 25 - unicasting said multicast packet to a corresponding well known end-node via a preestablished data channel; and
 - 30 - if said multicast group is reported to said well-known end-node, said well known end-node:

- 39 -

- extracting the multicast packet from the unicast packet; and
- sending the extracted multicast packet to group members using its point-multipoint tree to said members.

5

43. A method, as claimed in claim 42, characterised by a MCS wishing to participate in a multicast group, configuring two addresses, namely, ATM addresses of primary and secondary (back-up) well known end-nodes.

10

44. A method, as claimed in claim 43, characterised by a MCR maintaining:

- a single bi-directional-control connection to the well-known end-node for registration/de-registration purposes;
- a single data unidirectional data channel for tunneling multicast packets; and
- a number (N) of incoming multicast VCs, where N is the total number of multicast groups.

15

20

45. A method, as claimed in any of claims 42 to 44, characterised by:

- in the event of failure of a primary well known end-node, dismantling all point-to-point connections between the MCS routers and said primary well-known end-node;

25

- the MCRs, on detecting dismantlement of said connections, opening a new circuit to a back-up, pre-configured, well known end-node;

30

- each MCR registering the ATM addresses of the group members of which it is aware to the back-up well known end-node;
- said back-up well known end-node, on receipt of the ATM addresses,

- 40 -

establishing point-multipoint trees or each group;

- said MCRs continuously try to reestablish a connection with said primary well known end-node; and

5

- on reestablishment of said connection, said MCRs dismantle their connection to said back-up well known end-node and re-establish their connection to said primary well known end-node.

10

46. A telecommunication system, characterised in that said system includes an ATM transmission system as claimed in any one of the claims 1 to 20, or is adapted to operate in accordance with the method as claimed in any of claims 21 to 45.

15

20

25

30

1/3

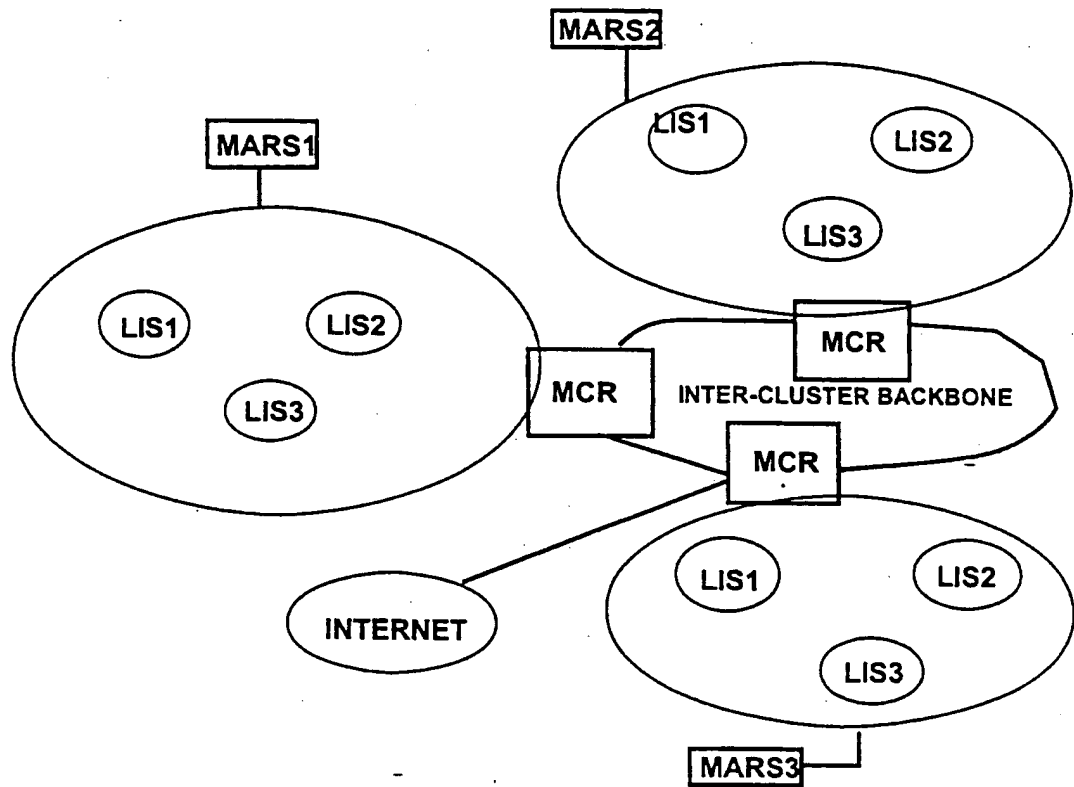


FIGURE 1




FIGURE 2

FIGURE 3




FIGURE 4




FIGURE 5